Constructing Performance Models for Dense Linear Algebra Algorithms on Cray XE Systems

Jorge González-Domínguez*[1], Evangelos Georganas†[1], Yili Zheng‡, María J. Martín*

*Computer Architecture Group, University of A Coruña, A Coruña, Spain
†Dept. of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720
‡Lawrence Berkeley National Laboratory, Berkeley, CA 94720

arXiv:1402.1285v1 [cs.DC] 6 Feb 2014

***Abstract*—Hiding or minimizing the communication cost is key in order to obtain good performance on large-scale systems. While communication overlapping attempts to hide communications cost, 2.5D communication avoiding algorithms improve performance scalability by reducing the volume of data transfers at the cost of extra memory usage. Both approaches can be used together or separately and the best choice depends on the machine, the algorithm and the problem size. Thus, the development of performance models is crucial to determine the best option for each scenario. In this paper, we present a methodology for constructing performance models for parallel numerical routines on Cray XE systems. Our models use portable benchmarks that measure computational cost and network characteristics, as well as performance degradation caused by simultaneous accesses to the network. We validate our methodology by constructing the performance models for the 2D and 2.5D approaches, with and without overlapping, of two matrix multiplication algorithms (Cannon's and SUMMA), triangular solve (TRSM) and Cholesky. We compare the estimations provided by these models with the experimental results using up to 24,576 cores of a Cray XE6 system and predict the performance of the algorithms on larger systems. Results prove that the estimations significantly improve when taking into account network contention.**

***Index Terms*—Performance Model, Performance Estimation, Communication Avoidance, Communication Overlapping, Network Contention**

I. Introduction

In spite of the improvement of the interconnection networks during recent years, communication cost is one of the most significant factors in application performance and the cost of data movement within and between nodes will continue to grow relative to the cost of computation. In fact, the effective bandwidth of many current networks decreases when many messages are concurrently sent from different cores.

Two techniques to minimize the impact of communication cost are *communication avoidance* and *communication overlapping*. The first one reduces the volume of communications [1], [2] while the latter hides communication impact on the application performance by overlapping communications with computational work. Nevertheless, complicated interactions and trade-offs arise when they are used together. Our previous work [3] studied how to combine and balance these two techniques for various linear algebra algorithms. Specifically, these techniques were applied to two matrix multiplication algorithms (Cannon's and SUMMA), triangular solve (TRSM) and Cholesky factorization. We demonstrated that the combination of both techniques is usually effective and significantly reduces execution time in most cases. However, depending on the algorithm, the problem size and the number of cores, using only overlapping can be more efficient.

The goal of this work is to create a methodology for constructing performance models that predict the effects of communication avoidance and overlapping on Cray XE systems. An important application of these performance models would be in guiding performance optimizations such as selecting the appropriate technique and tuning parameters. The new methodology includes a calibration factor to take into account performance degradation when concurrent communications from multiple nodes occur. It considers the performance of numerical computations, the latency and the ideal communication bandwidth of the network, as well as this calibration factor. These data are measured through completely portable benchmarks, so this methodology can be used for other architectures. The performance models obtained through the new methodology were evaluated on four algorithms (Cannon's, SUMMA, TRSM and Cholesky factorization), although they are easily adapted to other numerical algorithms. As shown by the experimental results, the inclusion of the calibration factor makes model predictions much more accurate.

The rest of this paper is organized as follows. Section II summarizes the related work. Section III presents the platform used for the experiments. Section IV describes the methodology followed to develop the models and the benchmarks employed. In Section V we explain in detail the models for two examples, the Cannon's and TRSM algorithms. Section VI compares the estimations of the models with the experimental results on Hopper, a Cray XE6 system, and gives performance predictions for our algorithms on larger systems. Finally, conclusions are discussed in Section VII.

II. Related Work

Communication avoidance and overlapping are techniques used to improve application performance on large supercomputers. Overlapping communication with computation or with other communication can increase efficiency without decreasing the communication volume. The benefits of overlapping

[1]Jorge and Evangelos contributed equally to this work.

can be very significant, for example, in one-sided communication [4], [5] and in collective communication [6]. On the other hand, the so-called 2.5D algorithms take advantage of extra memory to reduce the volume of communication along the critical path. The motivation of this type of communication avoiding algorithms and the explanation of why they attain the communication lower bounds can be found in [1] and [2]. These works also present the MPI implementation of some 2.5D algorithms and their experimental evaluation. As indicated in Section I, the interaction of both techniques for UPC-based implementations of Cannon's, SUMMA, TRSM and Cholesky algorithms was thoroughly evaluated in [3].

Regarding performance modeling, a well known communication model for distributed memory is the LogP family [7], [8], [9]. However, these models do not take into account the bandwidth degradation caused by the sharing of resources in large-scale systems, which has been proved as a crucial factor for communication performance on modern systems [10]. Several recent works began to take this issue into account to model current parallel applications. Hoefler and Snir [11] provided a definition of the terms *dilation* and *congestion* and considered them to generate efficient mapping policies for MPI applications. Bhathelé and Kalé [12] studied the effects of contention on large supercomputers with torus and mesh network topologies. They conclude that, in presence of contention, message latencies increase significantly with the distance between the communicated nodes. Performance models for interconnection networks that take congestion into consideration were presented in [13] and [14]. These models measure degradation of bandwidth using graphs and need a deep knowledge of the network schemes of communication.

Our model is based on benchmarking and can be considered as an adaptation to current systems of the $\alpha - \beta$ model used in [15] to estimate communication times. It takes into account the features of the modern machines by including factors that incorporate communication performance degradations.

There is also research done for shared memory platforms that consider contention as a key factor. For instance, Gibbons et al. [16] take into account memory contention while Helman and Jájá [17] also include processor contention. A two-level model that characterizes the impact of contention and cache effects and, at the same time, develops and studies a graph model of the application is presented in [18]. There exist similar works for hybrid MPI/OpenMP computation, that can help to decide the best combination of MPI processes and OpenMP threads depending on the characteristics of the machine and the code [19], [20], [21].

III. Experimental Platform

Before the explanation of the performance models, we describe the hardware and software environments used in the experiments because they are helpful to understand some of the results of the benchmarks.

Our target system is Hopper, a Cray XE6 supercomputer with 153,216 compute cores and 217 TB of memory in total. Each Cray XE6 node has 24 cores, grouped by 6 in 4 Non-Uniform Memory Access (NUMA) domains. CPU cores have faster access speed to memory within the same NUMA domain. Inter-node communication is done through the custom Cray Gemini Network, which is a high-bandwidth and low-latency 3D torus interconnect with hardware RDMA support. Table I lists the specifications of our experimental system.

System	Cray XE6 (Hopper)
Processor	AMD Opteron "Magny-Cours"
Clock rate	2.1 GHz
Peak performance per core	8.4 Gflops
Cores per NUMA domain	6
NUMA domains per node	4 (packaged in 2 sockets)
Total cores per node	24
Private L1 data cache per core	64 KB
Private L2 data cache per core	512 KB
Shared L3 cache per NUMA domain	6 MB
Memory bandwidth	25.6 GB/s
Memory per node	32 GB DDR3-1066 ECC
Compiler	Cray Compiler
Interconnect	Gemini 3D Torus
Peak Bandwidth (per direction)	7 GB/s

Table I
SPECIFICATIONS OF THE HOPPER SUPERCOMPUTER

The parallel implementations of the Cannon's, SUMMA, TRSM and Cholesky factorization algorithms presented in [3] are used as benchmark applications to test the methodology for constructing the proposed performance models. Although the codes are mainly written using UPC, the implementations employ three different parallel programming models (UPC, MPI and Pthreads) to exploit the hardware potential. These parallel algorithms use local linear algebra routines provided by the LibSci library, a collection of multithreaded numerical routines optimized for best performance on Cray systems. Thus, a hybrid (process/thread) parallelization model is used. Specifically, we run one UPC process per NUMA domain and one thread per core within the NUMA region. Each UPC process is mapped into an OS process and each process uses 6 OS threads, 1 thread per core. MPI collectives are used to supplement the UPC implementations in order to overcome some limitations of the current UPC collective library [3].

IV. Description of the Methodology

In this section we describe the general methodology to develop detailed performance models for linear algebra algorithms. They are constructed by using a divide-and-conquer methodology, which tracks the execution flow of each algorithm and estimates the completion time for every encountered operation. These operations are identified as computation or communication. In regard to the overlapped operations, the models predict the execution time as the maximum expected completion time of each individual operation. System parameters of the target computer extracted through portable benchmarks are used for the estimations.

Computation operations refer to the numerical computations that are performed by each process using its local data. Parallel algorithms usually employ multithreaded linear algebra routines provided by vendor optimized libraries for

local computation. Thus, their completion time depends on the computational efficiency of the routine on the system, the number of underlying threads and the problem size. Micro-benchmarks are executed on the target platform to obtain the efficiency of each employed routine. Functions $T_{rout}(d,t)$ return within the models the time to perform the computations of matrices with size d calling the $rout$ routine with t underlying threads. Both the benchmarks and the estimation functions are developed to predict the computational time of problems that work with square matrices. Operations with rectangular matrices are estimated as several consecutive square matrix operations.

Figure 1 illustrates the efficiency of the BLAS routines used by the benchmark applications on Hopper for different matrix sizes. As in the parallel benchmarks, six threads (one per core within the NUMA domain) were used to run the LibSci routines.

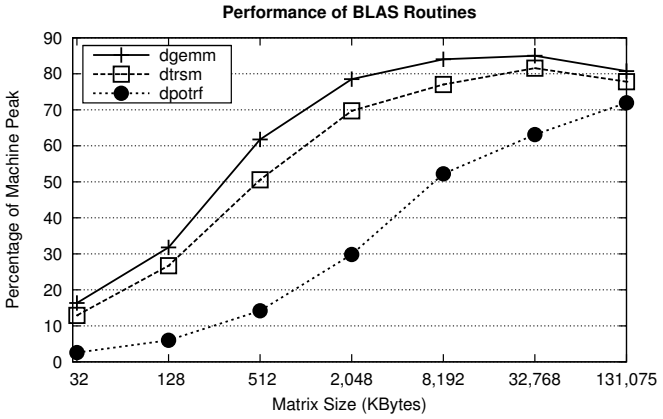

Figure 1. Efficiency of several LibSci BLAS routines on Hopper. Run with 6 cores in a NUMA domain.

The second type of operations is communication. In an ideal scenario, the time to transfer data between two processors should be easily determined by only knowing the network latency, the contention-free bandwidth between two nodes and the message size. If w the number of transferred elements, L the latency of the system, and β the inverse bandwidth (in seconds/word), we estimate itassume as:

$$T_{comm_ideal}(w) = L + \beta * w$$

Latency and bandwidth are obtained through the benchmarks presented for the LogP model [7]. It uses only two processes placed on different nodes and measures the communication times for different message sizes. Several executions are performed for each message size and the average communication times are used to calculate latency and bandwidth. Figure 2 shows the contention-free bandwidth of Hopper's network when UPC one-sided communication is used.

However, sharing the same network links among several nodes leads to a significant communication performance degradation in real environments where several parallel processes (created by the same or other programs) are transferring

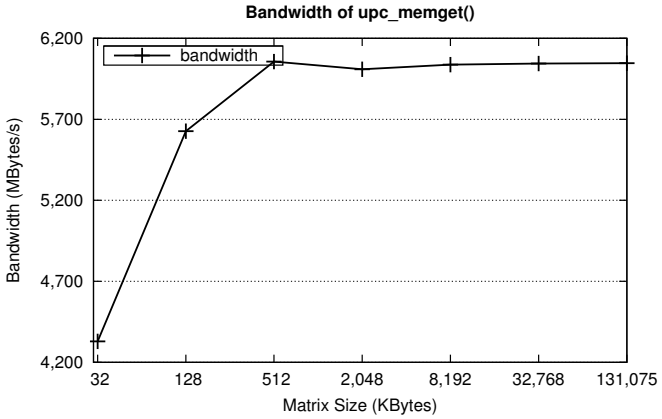

Figure 2. Inter-node contention-free network bandwidth on Hopper using UPC one-sided communication.

data simultaneously. A new micro-benchmark is therefore included to determine what we call "calibration factor", that is, the rate between ideal and real communication costs when several processes use the network simultaneously. The structure of this microbenchmark is the same as for the calculation of the ideal bandwidth in the LogP model but, in this case, several processes send the same message at the same time. Only one process per NUMA region is considered as multithreaded libraries are usually available and its use is recommended for the parallelization of linear algebra algorithms. The microbenchmark is executed varying the message size, the number of processes considered and the "communication distance" between processes. We define "communication distance" as the rank difference between the source process and the destination process. Although the mapping of processes to nodes is independent of the rank (processes with close ranks do not need to be scheduled on close nodes), generally larger distances imply messages moving along more links of the network.

For instance, Figure 3 shows the time needed by each process to perform 10 communication operations of 64MB, with a communication distance of 16 (64 processes were used for this experiment) and all the processes carrying out the communication at the same time. As can be seen in the figure, there is a huge variance among the times obtained by different processes. Degradation is not regular in all processes. Therefore, two calibration factors are extracted from each one of the micro-benchmark executions: the average and the maximum. Figure 4 shows these two calibration factors for 1,024 processes (6,144 cores) and 4,096 processes (24,576 cores) on Hopper. Messages of size 64MBytes were employed for this figure but additional experiments were performed changing the message size. The following conclusions were derived from these experiments:

- Both calibration factors do not significantly depend on the message size if messages above 256KBytes are used.
- The average calibration factor does not significantly depend on the total number of processes used at the same

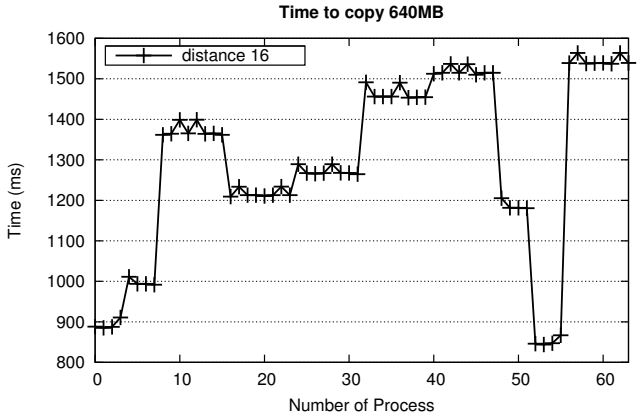

Figure 3. Completion time of each process to transfer 640MB using ten communications with a communication distance of 16 on Hopper (64 processes perform the same operation at the same time).

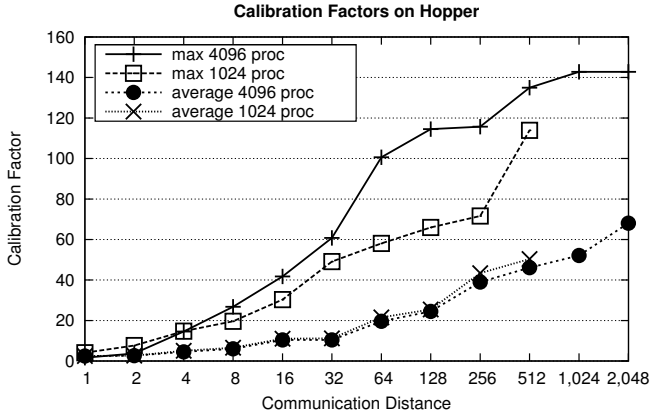

Figure 4. Calibration factors on Hopper depending on the distance of the processes that exchange the data.

time.

- The maximum calibration factor significantly depends on the total number of processes used at the same time.
- Both the average and the maximum calibration factors significantly depend on the communication distance. For instance, whatever the total number of processes is, the factor when all processes transfer data with a communication distance of 32 is much higher than the factor with a communication distance of 4. The increase of latency due to message distance on modern large-scale supercomputers with 3D networks had been previously analyzed in [12].

Therefore, the average calibration factor only depends on the communication pattern while the maximum also depends on the total number of processes performing a communication at the same time. Two functions were created in order to include these factors in the models:

- $C_{avg}(d)$: It returns the average calibration factor when all the participating processes perform communication at the same time with a communication distance d.
- $C_{max}(p, d)$: It provides the maximum calibration factor when p processes perform communication at the same time with a communication distance d.

In order to estimate the communication cost, we start executing the microbenchmark to determine the impact of network contention using a representative message size of 64MB and different numbers of processes and communication distances. Then, we obtain the calibration factors as the ratio between the ideal and the real communication time obtained on each scenario. For each number of threads we use a different execution on order to not leave nodes idle. These factors will be used into the models to better estimate the communication times. In presence of synchronizations all threads have to wait for the one with the highest communication cost, thus the maximum calibration factor will be used when synchronization among processes is required and the average will be used otherwise. Being p the total number of processes and d the communication distance, the completion time of each communication operation is estimated as follows (without and with synchronization, respectively):

$$T_{comm}(w, d) = C_{avg}(d) * (L + \beta * w)$$

$$T_{comm_sync}(p, w, d) = C_{max}(p, d) * (L + \beta * w)$$

V. Performance Models

The methodology proposed in the previous section was applied to model the implementations of Cannon's, SUMMA, TRSM and Cholesky factorization presented in [3]. In order to simplify the paper, only the Cannon's and TRSM algorithms are discussed in detail as representative examples.

A. Cannon's Algorithm

The traditional 2D Cannon's algorithm [22] performs the matrix-matrix multiplication ($A * B = C$) by shifting blocks of data among near neighbor processes on a 2D grid of p processes, with square blocks of size $n/\sqrt{p}$ (being n the total matrix size). The algorithm starts by performing a skew on the initial matrices along rows and columns of the process grid. The blocks are lined up so that at each subsequent step only a single shift needs to be done between each block multiplication. Cannon's algorithm uses near-neighbor point-to-point remote copies rather than collective communications.

The 2D Cannon's algorithm presents the following structure:

1) An initial shift where each process copies a block of matrix A from one process that is in its row of the grid and one block of matrix B from one process that is in its column of the grid.
2) A loop with $\sqrt{p} - 1$ iterations where:
   a) Each process calls *dgemm* with the current blocks.
   b) The processes in the same rows and in the same columns of the grid are synchronized to guarantee

that all of them have received the blocks that will be copied in the next iteration.

c) Each process copies the next blocks. The block of A is copied from the next process in the row ($MYTHREAD$[2]$+1$) and the block of B from the next process in the column ($MYTHREAD+\sqrt{p}$).

3) A *dgemm* operation with the final blocks.

As explained in Section IV the estimation of the execution time for the whole algorithm is the addition of the estimation of all parts. Thus, it can be specified as:

$$T_{Cannon_2D} = T_{iniShiftRow} + T_{iniShiftCol} + (\sqrt{p}-1) * \Big(T_{shiftRow} + T_{shiftCol} + T_{multBlock}\Big) + T_{multFinalBlock}$$

This equation is easily simplified because the initial shifts and the final *dgemm* are equal to the shifts and products of the loop, respectively. As explained in Section IV, T_{dgemm}, obtained from the execution of the multithreaded $dgemm$ BLAS routine on the target platform, provides an estimation of the computational time of this routine. Regarding the communications, some considerations must be taken into account:

- Shifts among processes in the same column of the grid involve a communication of distance $\sqrt{p}$, whereas shifts among processes in the same row of the grid involve a communication of distance 1.
- Due to the synchronizations within the loop, all processes must wait for the last process that finishes the shifts. The effect of this synchronization is included in the model by using the maximum calibration factor.

With all these assumptions the completion time of each communication operation (including the impact of the synchronization) is estimated using T_{comm_sync} (see Section IV). Consequently, as the communication distances in the shifts by rows and columns are, respectively, 1 and $\sqrt{p}$, the model for the 2D Cannon's algorithm is:

$$T_{Cannon_2D} = \sqrt{p} * \Big(T_{comm_sync}(p, bs^2, 1) + T_{comm_sync}(p, bs^2, \sqrt{p}) + T_{dgemm}(bs, t)\Big)$$

bs being the block size ($n/\sqrt{p}$) and t the number of underlying threads used to run the local numerical routine.

Overlapping is included in this algorithm by moving forward the shifts of the next iteration of the loop and hiding their cost by using asynchronous communications while performing the products with the blocks of the current iteration. We consider that the algorithm achieves perfect overlapping, so the completion time is estimated as the maximum of the two operations. The model is the addition of the completion time of the first shift, the final *dgemm* (which cannot be overlapped) and the overlapped loop:

$$T_{Cannon_2D_ovlp} = T_{comm_sync}(p, bs^2, 1) + T_{comm_sync}(p, bs^2, \sqrt{p}) + T_{dgemm}(bs, t) + (\sqrt{p}-1) * max\Big[T_{comm_sync}(p, bs^2, 1) + T_{comm_sync}(p, bs^2, \sqrt{p}), T_{dgemm}(bs, t)\Big]$$

The 2.5D algorithm initially replicates the blocks of A and B on c layers and redundantly performs independent updates on the copies of C, which are combined using a reduction operation at the end. The blocks size is $n/(\sqrt{p}/c)$ but there are only $\sqrt{p}/c$ shifts, decreasing the total amount of communications. The initial replications are performed using communications whose source process is always on the first layer. The communication distance depends on the target layer. The worst case (distance to the last layer) will be used for the estimation:

$$T_{iniRepl}(p, w, c) = 2 * C_{max}(p, (c-1) * p/c) * (L + \beta * w)$$

The final reduction is performed using the MPI collective reduce operation. The algorithms used for collective communications depend on the MPI implementation. We assume that the behavior of collective operations is the one specified in [23] where the authors study different algorithms to perform these operations and indicate which is the best option depending on the number of processes and the message size. Taking into account that power-of-two numbers of processes and medium-sized messages are used by our benchmark applications, the Rabenseifner's algorithm [24], a reduce-scatter followed by a gather to the root with a synchronization between them, is assumed for the reduce operation. Considering the parameter d as the closest communication distance between two of the processes involved in the reduce, q the number of processes involved in the reduction, p the total number of processes, and w the size of the block to reduce, the estimation function is:

$$T_{reduce}(p, q, w, d) = T_{redSca_sync}(p, q, w, d) + T_{gather}(q, w, d)$$

The reduce-scatter follows the recursive halving algorithm (see [23]). If there are q processes involved in the reduction, there are $\log_2(q)$ steps where each process exchanges data with a process that is in a distance of 2^i (i being the number of the step). Consequently, each step increases the communication distance. In order to include the impact of the synchronization between the reduce-scatter and the gather, the last step uses the maximum calibration factor. For the rest of the steps the average factor is used:

$$T_{redSca_sync}(p, q, w, d) = \sum_{i=0}^{\log_2(q)-2} \Big(C_{avg}(2^i * d) * (L + \beta * w * q/2^i)\Big) + C_{max}(p, 2^{\log_2(q)-1} * d) * (L + \beta * w * t/2^{\log_2(q)-1})$$

[2] $MYTHREAD$ is the UPC indentifier for the thread that is performing that part of the code

The binomial tree algorithm is used for the gather. The average calibration factor is always employed in the equation as there is no synchronization at the end:

$$T_{gather}(q,w,d) = \sum_{i=0}^{\log_2(q)-1} \Big(C_{avg}(2^i * d) * (L + \beta * (w/q) * 2^i)\Big)$$

Therefore, knowing that c is the number of layers and thus the block size is $n/\sqrt{p/c}$, the performance models for the 2.5D algorithms are:

$$\begin{aligned} T_{Cannon_2.5D} = {} & T_{iniRepl}(p, bs^2, c) + (\sqrt{p/c} - 1) * \\ & \Big(T_{comm}(p, bs^2, 1) + T_{comm}(p, bs^2, \sqrt{p/c}) + \\ & T_{dgemm}(bs,t)\Big) + T_{dgemm}(bs,t) + \\ & T_{reduce}(p, c, bs^2, p/c) \end{aligned}$$

$$\begin{aligned} T_{Cannon_2.5D_ovlp} = {} & T_{iniRepl}(p, bs^2, c) + (\sqrt{p/c} - 1) * \\ & max\Big[T_{comm}(p, bs^2, 1) + \\ & T_{comm}(p, bs^2, \sqrt{p/c}), T_{dgemm}(bs,t)\Big] + \\ & T_{dgemm}(bs,t) + T_{reduce}(p, c, bs^2, p/c) \end{aligned}$$

B. Triangular Solve

Triangular solve (TRSM) is used to compute a matrix X, such that $X * U = B$, where U is upper-triangular and B is a dense matrix. This problem has dependencies across the columns of X, while each row of X can be computed independently.

The 2D algorithm is based on the block version where the matrices are partitioned in blocks as:

$$\begin{bmatrix} X_{00} & X_{01} \\ X_{10} & X_{11} \end{bmatrix} * \begin{bmatrix} U_{00} & U_{01} \\ 0 & U_{11} \end{bmatrix} = \begin{bmatrix} B_{00} & B_{01} \\ B_{10} & B_{11} \end{bmatrix}$$

and, at each step, a block-column of X is computed. This block-column can then be used to update the trailing matrix B. The computation proceeds as follows:

1) Broadcast U_{00} and U_{01} along the columns of the grid.
2) Compute via *dtrsm* $X_{00} = B_{00} * U_{00}^{-1}$ and $X_{10} = B_{10} * U_{00}^{-1}$.
3) Once the previous *dtrsm* are completed, broadcast X_{00} and X_{10} along the rows of the grid.
4) Update via *dgemm* $B_{01} = B_{01} - X_{00} * U_{01}$ and $B_{11} = B_{11} - X_{10} * U_{01}$.
5) Broadcast U_{11} along the columns of the grid.
6) Compute via *dtrsm* $X_{01} = B_{01} * U_{11}^{-1}$ and $X_{11} = B_{11} * U_{11}^{-1}$.

Communication is minimized by employing a block-cyclic layout where each process owns multiple sub-blocks of the matrices. The number of blocks per process is indicated by the parameter r. Consequently, the matrices are divided in $r*\sqrt{p}$-by-$r*\sqrt{p}$ blocks and the block size is $bs = n/(r*\sqrt{p})$. Thus, steps 1-4 are computed inside a loop of $r * \sqrt{p}$ iterations.

The main part of the algorithm is a loop with four operations per iteration. The estimation is the addition of the completion time of each iteration plus the final broadcast and *dtrsm*:

$$\begin{aligned} T_{TRSM_2D} = {} & \sum_{i=0}^{r*\sqrt{p}-1} \Big(T_{bcast_U} + T_{solve_iter} + T_{bcast_X} + \\ & T_{update}\Big) + T_{last_bcast_U} + T_{last_solve} \end{aligned}$$

The main difference with the matrix-matrix multiplication algorithms is that the workload changes according to the iteration. The more iterations pass, the less blocks of U are involved in the update and the broadcast along the columns of the grid. Besides, most of computational operations work with rectangular matrices which are estimated as several calls using square matrices, as explained in Section IV.

The broadcasts and gathers in the different TRSM versions are performed employing MPI. Following again the work [23], the MPI broadcast is assumed to be performed through a scatter followed by an all-gather (with a synchronization between them):

$$T_{bcast}(p,q,w,d) = T_{scatter_sync}(p,q,w,d) + T_{all-gather}(q,w,d)$$

The models for $T_{scatter_sync}$ and $T_{all-gather}$ are exactly the same as the ones for T_{redSca_sync} and T_{gather} presented for Cannon's algorithm, respectively. Finally, it must be taken into account that there is one synchronization per iteration: the broadcast of X along rows cannot start until the previous computations have finished. This synchronization is modeled by using T_{bcast_sync} for the broadcast of U. It uses the maximum calibration factor in the last of the $\log_2(q)$ steps of the gather. These assumptions lead to the following model:

$$\begin{aligned} T_{TRSM_2D} = {} & \sum_{i=0}^{r*\sqrt{p}-1} \Big(((r*\sqrt{p} - i)/\sqrt{p}) * \\ & T_{bcast_sync}(p, \sqrt{p}, bs^2, \sqrt{p}) + r * (T_{dtrsm}(bs,t) + \\ & T_{bcast}(p, \sqrt{p}, bs^2, 1)) + ((r*\sqrt{p} - i - 1)/\sqrt{p}) * \\ & T_{dgemm}(bs^2, t))\Big) + r * T_{dtrsm}(bs,t) + \\ & T_{bcast_sync}(p, \sqrt{p}, bs^2, \sqrt{p}) \end{aligned}$$

The algorithm is optimized by overlapping the broadcast of the next iteration of U with the matrix update. In this case, Pthreads are used so that one of the threads is completely dedicated to communications. Therefore, these communications are overlapped with the numerical computation performed by the other $t-1$ threads.

$$\begin{aligned} T_{TRSM_2D_ovl} = {} & r * T_{bcast_sync}(p, \sqrt{p}, bs^2, \sqrt{p}) + \\ & \sum_{i=0}^{r*\sqrt{p}-1} \Big(r * (T_{dtrsm}(bs, t-1) + \\ & T_{bcst}(p, \sqrt{p}, bs^2, 1)) + ((r*\sqrt{p} - i - 1)/\sqrt{p}) * \\ & max[T_{bcast_sync}(p, \sqrt{p}, bs^2, \sqrt{p}), \\ & r * T_{dgem}(bs^2, t-1)]\Big) + r * T_{dtrsm}(bs, t-1) \end{aligned}$$

For the 2.5D approach, the initial distribution of the matrices in one layer is the same as for the 2D algorithm: a block-cyclic distribution along the p/c processes in each dimension. Thus, the block size is $n/(r*\sqrt{p/c})$-by-$n/(r*\sqrt{p/c})$. We initially replicate the triangular matrix U along layers (as the input matrices in the 2.5D algorithm for the matrix product) but distribute the rows of each block of X among them using one scatter operation. Remark that there are two levels of blocking for X, thus each process has r^2 rectangular blocks of $n/(c*r*\sqrt{p/c})$-by-$n/(r*\sqrt{p/c})$ elements. With this distribution, each layer computes a subset of the rows of X via a 2D TRSM with its $\sqrt{p/c}$-by-$\sqrt{p/c}$ grid of processes. The distributed rows must be gathered once the layers finish their computation. The models for this 2.5D algorithm are:

$$
\begin{aligned}
T_{TRSM_2.5D} = r^2 * \Big(& (3/4) * T_{bcast}(p, c, bs^2, p/c) + \\
& T_{scatter_sync}(p, c, bs^2/c, p/c) \Big) + \\
& \sum_{i=0}^{r*\sqrt{p/c}-1} \Big(((r * \sqrt{p/c} - i)/\sqrt{p/c}) * \\
& T_{bcast_sync}(p, \sqrt{p}, bs^2, \sqrt{p/c}) + (r/c) * \\
& (T_{dtrsm}(bs, t) + T_{bcast}(p, \sqrt{p/c}, bs^2, 1) + \\
& ((r * \sqrt{p/c} - i - 1)/\sqrt{p/c}) * \\
& T_{dgemm}(bs^2, t)) \Big) + \\
& T_{bcast_sync}(p, \sqrt{p/c}, bs^2, \sqrt{p/c}) + (r/c) * \\
& T_{dtrsm}(bs, t) + r^2 * T_{gather}(c, bs^2, p/c)
\end{aligned}
$$

$$
\begin{aligned}
T_{TRSM_2.5D_ovlp} = r^2 * \Big(& (3/4) * T_{bcast}(p, c, bs^2, p/c) + \\
& T_{scatter_sync}(p, c, bs^2/c, p/c) \Big) + \\
& r * T_{bcast_sync}(p, \sqrt{p/c}, bs^2, \sqrt{p/c}) + \\
& \sum_{i=0}^{r*\sqrt{p/c}-1} \Big((r/c) * (T_{dtrsm}(bs, t-1) + \\
& T_{bcast}(p, \sqrt{p/c}, bs^2, 1)) + \\
& ((r * \sqrt{p/c} - i - 1)/\sqrt{p/c}) * \\
& max[T_{bcast_sync}(p, \sqrt{p/c}, bs^2, \sqrt{p/c}), \\
& (r/c) * T_{dgemm}(bs^2, t-1)] \Big) + \\
& (r/c) * T_{dtrsm}(bs, t-1) + \\
& r^2 * T_{gather}(c, bs^2, p/c)
\end{aligned}
$$

VI. Experimental Evaluation

This section evaluates the performance models for Cannon's, SUMMA, TRSM and Cholesky factorization using the methodology explained in this work. Hopper, the supercomputer presented in Section III, was used as testbed. We study not only the accuracy of the estimations provided by the models compared with the real performance, but also the predictions that they provide for scenarios with more resources.

A. Validation of the Proposed Models

Experimental results of the benchmarking algorithms were obtained on Hopper in order to compare them with the estimations provided by the models. The graphs presented in this section use the percentage of peak flops of the target system as performance measure. The comparison is shown for 6,144 and 24,576 cores (1,024 and 4,096 processes). In order to illustrate that the inclusion of the calibration factor is key to obtain good estimations both results with and without considering the calibration factor are included in the figure (`est_Cal` and `est_NoCal`, respectively).

Figures 5 and 6 illustrate the comparison for the matrix multiplication algorithms (Cannon's and SUMMA, respectively) using matrices with 32,768×32,768 doubles. It must be remarked that the estimations using the calibration factor allow to rank the different versions of both algorithms correctly according to their performance. The differences with the real results are always less than 4% of the machine peak for Cannon's and 7% for SUMMA.

The comparison of the experimental results and estimations for TRSM and Cholesky factorization with matrices of 65,536×65,536 doubles are shown in Figures 7 and 8, respectively. They demonstrate that the increase of complexity of the algorithm does not lead to worse estimations.

B. Predictions for Executions with Additional Resources

As previously mentioned, one of the main advantages of having performance models available is to assist deciding which is the best algorithm for a certain scenario (type of machine, number of cores and problem size). Furthermore, programmers can also benefit from performance models to simulate what would happen on scenarios with more resources than the available testbeds. Which algorithm would be the best if infinite memory was available? What would be the trend of the algorithms if the number of cores was increased? These questions can be easily answered with appropriate performance models.

Tables II - V summarize the performance estimations for the four approaches of Cannon's, SUMMA, TRSM and Cholesky factorization. The approach with the best performance for each scenario (number of cores and matrix size) is in boldface. Estimations are obtained for matrices larger than those we could experimentally test and up to 393,216 cores. The maximum calibration factor is the only parameter used in our methodology that depends on the total number of cores. We could not execute the benchmark to measure this factor for more than 24,576 cores. A polynomial regression was employed to estimate the maximum calibration factor for a larger number of cores.

The behavior of the two matrix multiplication algorithms (Cannon's and SUMMA) is similar. If the workload per core is quite large the 2D approach with overlapping obtains the highest performance. Nevertheless, when the number of cores increases without changing the matrix size the avoidance of communication is more effective. Thus, for a certain matrix size there is a point (in terms of number of cores) where the

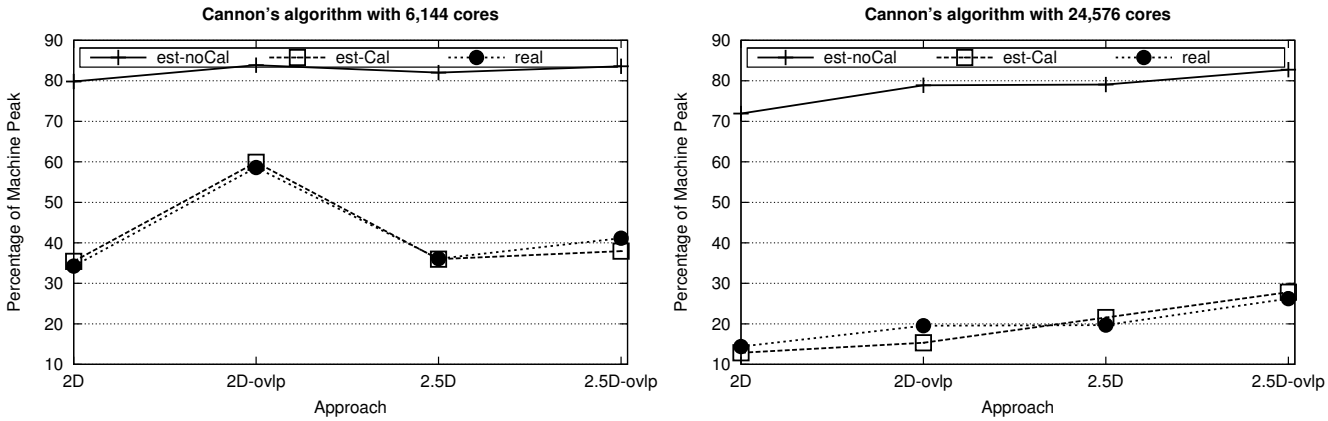

Figure 5. Real and estimated performance for Cannon's using matrices of 32,768×32,768 doubles

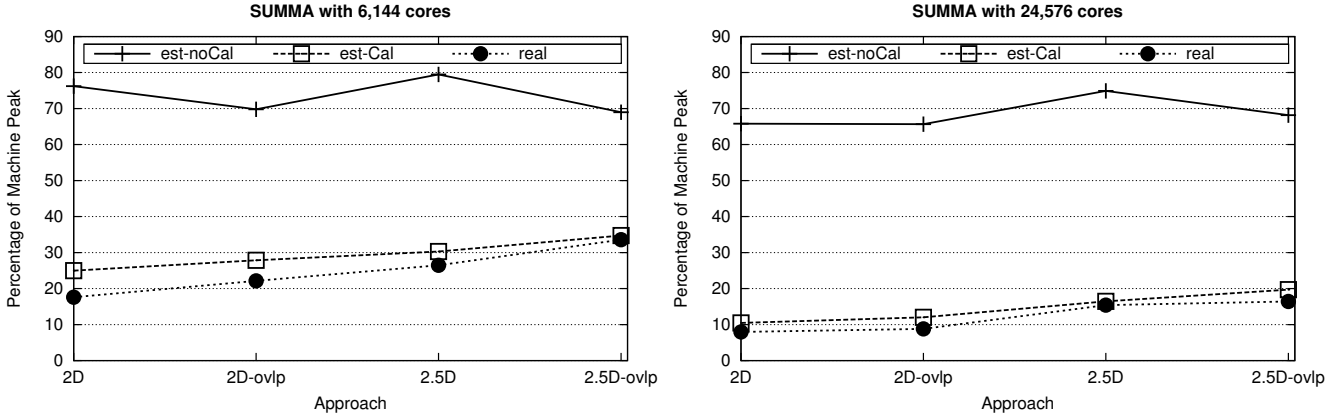

Figure 6. Real and estimated performance for SUMMA using matrices of 32,768×32,768 doubles

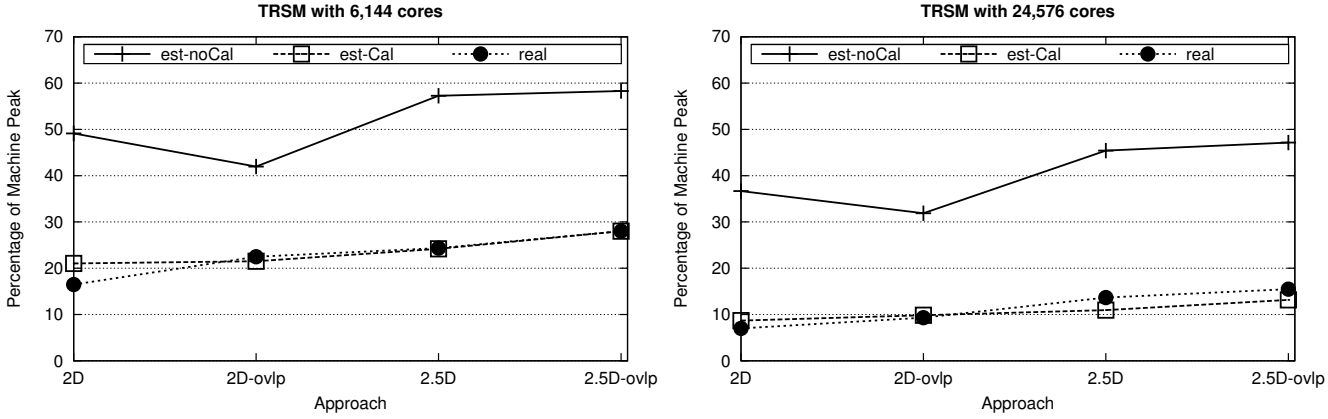

Figure 7. Real and estimated performance for TRSM using matrices of 65,536×65,536 doubles

2.5D algorithm with overlapping becomes the best choice. This behavior is not reproduced in the TRSM case, where the model predicts that the 2.5D algorithm with overlapping is the best option in all cases. Regarding Cholesky, as in the matrix multiplication algorithms, there is a sweet-spot in terms of number of cores where the 2.5D version with overlapping outperforms the other ones. These results confirm that communication avoiding algorithms combined with overlapping are needed on the way to the exascale era.

VII. Conclusions

In this paper we described how to construct performance models for parallel linear algebra algorithms that use communication avoiding and communication overlapping techniques.

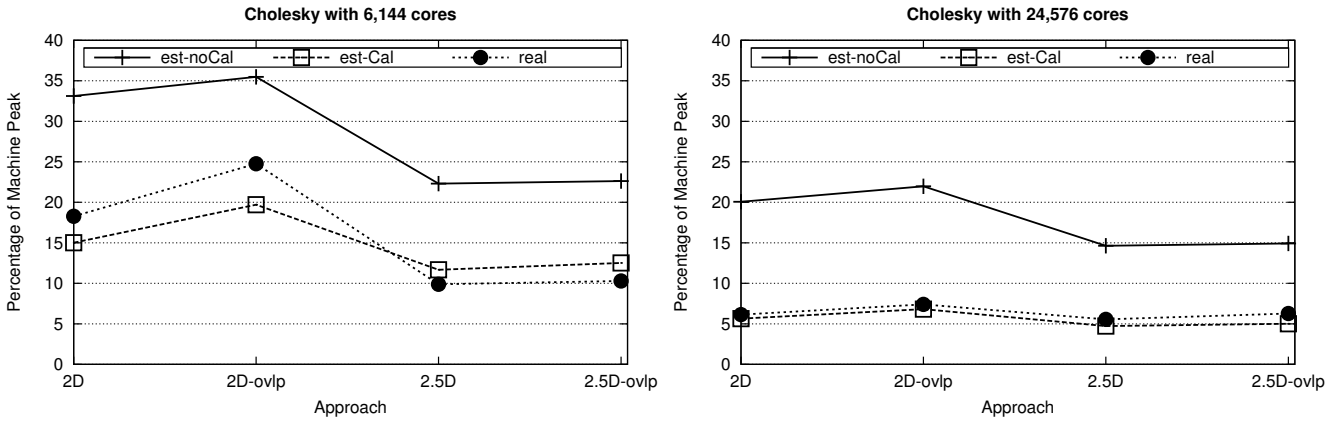

Figure 8. Real and estimated performance for Cholesky using matrices of 65,536×65,536 doubles

Size →	32768				65536			
Cores ↓	2D	2D_overlp	2.5D	2.5D_overlp	2D	2D_overlp	2.5D	2.5D_overlp
1,536	67.95	**83.69**	53.63	55.56	72.36	**80.40**	64.52	65.91
6,144	35.42	**59.88**	35.95	37.96	50.20	**73.20**	48.22	50.95
24,576	12.87	15.33	21.56	**27.80**	22.59	30.73	34.51	**45.78**
98,304	4.57	4.93	9.37	**10.55**	8.71	9.78	17.04	**21.04**
393,216	1.30	1.35	3.94	**4.19**	2.78	2.91	7.55	**8.32**

Table II
PREDICTED PERCENTAGE OF PEAK FLOP FOR THE CANNON'S ALGORITHM.

Size →	32768				65536			
Cores ↓	2D	2D_overlp	2.5D	2.5D_overlp	2D	2D_overlp	2.5D	2.5D_overlp
1,536	52.29	**68.59**	49.18	46.65	62.43	**66.47**	61.19	55.19
6,144	24.98	27.85	30.28	**34.74**	38.82	**58.69**	43.54	43.37
24,576	10.46	12.02	16.44	**19.71**	18.92	24.28	27.67	**38.51**
98,304	4.01	4.29	7.93	**8.75**	8.75	9.83	14.68	**17.51**
393,216	1.27	1.33	3.56	**3.77**	3.62	3.84	7.75	**8.56**

Table III
PREDICTED PERCENTAGE OF PEAK FLOP FOR THE SUMMA ALGORITHM.

Size →	65536				131072			
Cores ↓	2D	2D_overlp	2.5D	2.5D_overlp	2D	2D_overlp	2.5D	2.5D_overlp
1,536	43.40	39.85	41.37	**44.16**	56.10	49.62	55.58	**57.89**
6,144	21.04	21.50	24.20	**28.00**	33.49	32.39	38.01	**42.03**
24,576	8.70	9.84	10.94	**13.16**	15.87	17.10	20.12	**26.06**
98,304	3.33	3.60	4.42	**4.79**	6.85	7.88	9.13	**10.59**
393,216	1.24	1.29	1.38	**1.43**	2.87	3.06	3.11	**3.29**

Table IV
PREDICTED PERCENTAGE OF PEAK FLOP FOR THE TRIANGULAR SOLVE (TRSM).

Size →	65536				131072			
Cores ↓	2D	2D_overlp	2.5D	2.5D_overlp	2D	2D_overlp	2.5D	2.5D_overlp
1,536	32.29	**32.29**	21.02	21.81	46.88	**58.26**	29.86	30.72
6,144	15.02	**19.71**	11.68	12.51	18.44	**26.19**	14.78	15.96
24,576	5.64	**6.82**	4.73	5.01	6.36	**8.79**	6.47	6.60
98,304	1.89	**2.01**	1.83	1.87	4.67	**5.45**	4.29	4.29
393,216	0.56	0.57	0.59	**0.61**	1.66	1.74	1.76	**1.83**

Table V
PREDICTED PERCENTAGE OF PEAK FLOP FOR CHOLESKY.

The proposed methodology was evaluated using two matrix multiplication algorithms (Cannon's and SUMMA), triangular solve (TRSM) and Cholesky factorization. For each of them, we explored different optimization techniques and constructed

performance models for every version: communication avoiding, communication overlapping, and both.

The proposed models rely on the use of different parameters measured through portable benchmarks. These parameters are the efficiency of the computational kernels involved in the linear algebra algorithms (calls to multithreaded BLAS routines in our case), the latency and the ideal bandwidth of the network, and calibration factors that encapsulate the communication performance degradation that occurs when several processes concurrently access the network.

For all algorithms in this paper, the performance models can accurately predict the empirical results from the experiments on our target machine, a Cray XE6 system. We have also confirmed that incorporating the communication calibration factor in the models is critical to obtain correct estimations. Moreover, our models are quite flexible because they can take into account runtime constraints (e.g., available memory). These performance models can be useful in a number of applications, such as performance tuning for current systems, performance extrapolation for future systems, and new algorithm design.

Finally, although the methodology was developed focusing on a specific target architecture, it can be exported to systems with other characteristics as the benchmarks are portable. Proving if this methodology is enough to characterize the performance of dense linear algebra algorithms on other types of large supercomputers or if more parameters must be taken into account is considered future work.

ACKNOWLEDGMENTS

This research was supported in part by the Office of Science of the U.S. Department of Energy (DE-AC02-05CH11231), DARPA (HR0011-10-9-0008), the Ministry of Science and Innovation of Spain (TIN2010-16735) with FEDER funds of the European Union and the Ministry of Education of Spain under the FPU research grant AP2008-01578. This research used resources of the National Energy Research Scientific Computing Center. We gratefully thank Costin Iancu and Katherine Yelick for their valuable comments to improve this work.

REFERENCES

[1] E. Solomonik and J. Demmel, "Communication-Optimal Parallel 2.5D Matrix Multiplication and LU Factorization Algorithms," in *Proc. 17th Intl. European Conf. on Parallel and Distributed Computing (Euro-Par'11)*, ser. Lecture Notes in Computer Science, vol. 6853, Bordeaux, France, 2011, pp. 90–109.

[2] ——, "2.5D Algorithms for Parallel Dense Linear Algebra," *Concurrency and Computation: Practice and Experience (ToAppear)*, To Appear.

[3] E. Georganas, J. González-Domínguez, E. Solomonik, Y. Zheng, J. Touriño, and K. Yelick, "Communication Avoiding and Overlapping for Numerical Linear Algebra," in *Proc. 24th ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC'12)*, Salt Lake City, UT, USA, 2012.

[4] C. Bell, D. Bonachaea, R. Nishtala, and K. Yelick, "Optimizing Bandwidth Limited Problems Using One-Sided Communication and Overlap," in *Proc. 20th Intl. Parallel and Distributed Processing Symp. (IPDPS'06)*, Rhodes Island, Greece, 2006.

[5] R. Nishtala, P. Hargrove, D. Bonachea, and K. Yelick, "Scaling Communication-Intensive Applications on BlueGene/P Using One-Sided Communication and Overlap," in *Proc. 23rd Intl. Parallel and Distributed Processing Symp. (IPDPS'09)*, Rome, Italy, 2009.

[6] R. Nishtala, Y. Zheng, P. Hargrove, and K. Yelick, "Tuning Collective Communication for Partitioned Global Address Space Programming Models," *Parallel Computing*, vol. 37, no. 9, pp. 576–591, 2011.

[7] D. E. Culler, R. M. Karp, D. A. Patterson, A. Sahay, K. E. Schauser, E. Santos, R. Subramonian, and T. von Eicken, "LogP: Towards a Realistic Model of Parallel Computation," in *Proc. 4th ACM SIGPLAN Symp. on Principles and Practice of Parallel Programming (PPoPP'93)*, San Diego, CA, USA, 1993, pp. 1–12.

[8] A. Alexandrov, M. F. Ionescu, K. E. Schauser, and C. Scheiman, "LogGP: Incorporating Long Messages into the LogP Model," *Journal of Parallel and Distributed Computing*, vol. 44, no. 1, pp. 71–79, 1997.

[9] L. Yuan, Y. Zhang, Y. Tang, L. Rao, and X. Sun, "LogGPH: A Parallel Computational Model with Hierarchical Communication Awareness," in *Proc. 13th IEEE Intl. Conf. on Computational Science and Engineering (CSE'10)*, Hong Kong, China, 2010, pp. 268–264.

[10] A. Hoisie, G. Johnson, D. J. Kerbyson, M. Lang, and S. Pakin, "A Performance Comparison Through Benchmarking and Modeling of Three Leading Supercomputers: Blue Gene/L, Red Storm, and Purple," in *Proc. 18th ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC'06)*, Tampa Bay, FL, USA, 2006.

[11] T. Hoefler and M. Snir, "Generic Topology Mapping Strategies for Large-Scale Parallel Architectures," in *Proc. 25th Intl. Conf. on Supercomputing (ICS'11)*, Tucson, AZ, USA, 2011, pp. 75–84.

[12] A. Bhatelé and L. V. Kalé, "An Evaluative Study on the Effect of Contention on Message Latencies in Large Supercomputers," in *Proc. 23rd Intl. Parallel and Distributed Processing Symp. (IPDPS'09)*, Rome, Italy, 2009.

[13] J. Vienne, M. Martinasso, J.-M. Vincent, and J.-F. Méhaut, "Predictive Models for Bandwidth Sharing in High Performance Clusters," in *Proc. 10th IEEE Intl. Conf. on Cluster Computing (CLUSTER'08)*, Tsukuba, Japan, 2008, pp. 286–291.

[14] M. Martinasso and J.-F. Méhaut, "A Contention-Aware Performance Model for HPC-Based Networks: A Case Study of the InfiniBand Network," in *Proc. 17th Intl. European Conf. on Parallel and Distributed Computing (Euro-Par'11)*, ser. Lecture Notes in Computer Science, vol. 6852, Bordeaux, France, 2011, pp. 91–102.

[15] R. W. Hockney, "The Communication Challenge for MPP: Intel Paragon and Meiko CS-2," *Parallel Computing*, vol. 20, no. 3, pp. 389–398, 1994.

[16] P. B. Gibbons, Y. Matias, and V. Ramachandran, "Can Shared-Memory Model Serve as a Bridging Model for Parallel Computation?" in *Proc. 9th Annual ACM Symp. on Parallel Algorithms and Architectures (SPAA'09)*, New Port, RI, USA, 1997, pp. 72–83.

[17] D. R. Helman and J. Jájá, "Prefix Computations on Symmetric Multiprocessors," *Journal of Parallel and Distributed Computing*, vol. 61, no. 2, pp. 267–278, 2001.

[18] V. S. Adve and M. K. Vernon, "Parallel Program Performance Prediction Using Deterministic Task Graph Analysis," *ACM Transactions on Computer Systems*, vol. 22, no. 1, pp. 94–136, 2004.

[19] L. Adhianto and B. Chapman, "Performance Modeling of Communication and Computation in Hybrid MPI and OpenMP Applications," *Simulation Modelling Practice and Theory*, vol. 15, no. 4, pp. 481–491, 2007.

[20] L. M. Liebrock and S. P. Goudy, "Methodology for Modelling SPMD Hybrid Parallel Computation," *Concurrency and Computation: Practice and Experience*, vol. 20, no. 8, pp. 903–940, 2008.

[21] X. Wu and V. Taylor, "Performance Modeling of Hybrid MPI/OpenMP Scientific Applications on Large-scale Multicore Cluster Systems," in *Proc. 14th IEEE Intl. Conf. on Computational Science and Engineering (CSE'11)*, Dalian, China, 2011, pp. 181–190.

[22] L. E. Cannon, "A Cellular Computer to Implement the Kalman Filter Algorithm," Ph.D. dissertation, Montana State University, 1969.

[23] R. Thakur, R. Rabenseifner, and W. Gropp, "Optimization of Collective Communication Operations in MPICH," *Intl. Journal of High Performance Computing Applications*, vol. 19, no. 1, pp. 49–66, 2005.

[24] R. Rabenseifner, "Optimization of Collective Reduction Operations," in *Proc. 3rd Intl. Conf. on Computational Science (ICCS'04)*, ser. Lecture Notes in Computer Science, vol. 3036, Krakow, Poland, 2004, pp. 1–9.